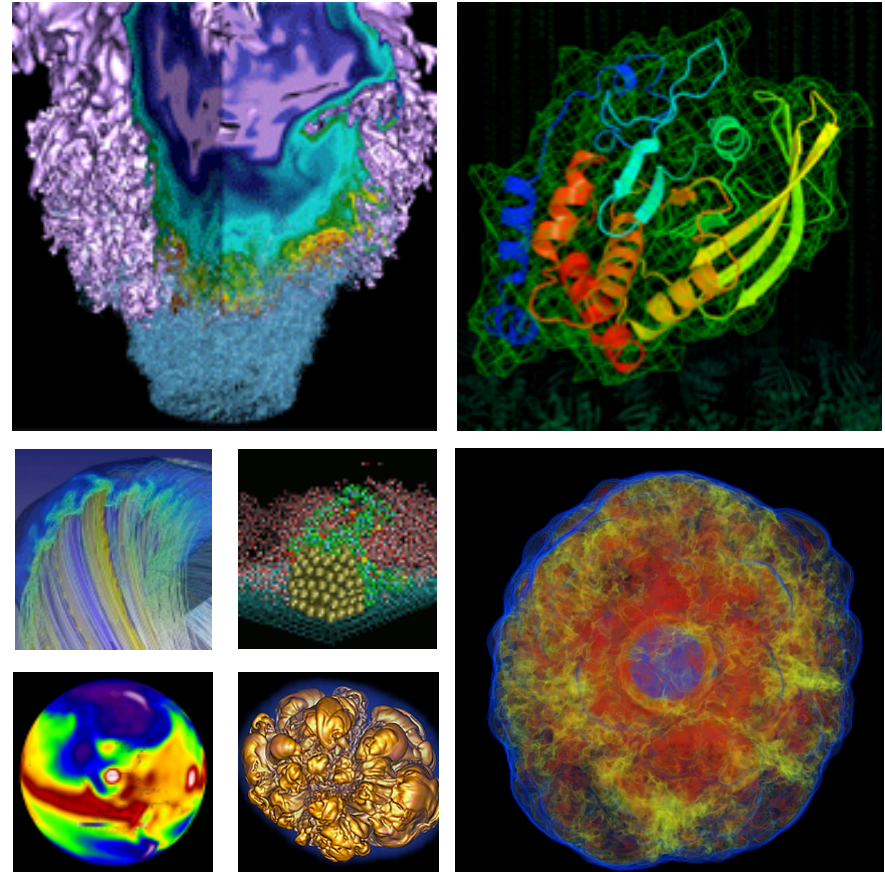# NERSC User Environment

David Turner
NERSC User Services Group

September 10, 2013

Overview

- **Login Nodes, File Systems, and Dot Files**
  - David Turner

- **Software Modules**
  - Doug Jacobsen

- **Compilers**
  - Mike Stewart

Login Nodes

- **Edison**
  - Six nodes
    - 16 cores, 2.0GHz Intel Sandy Bridge
    - 512GB

- **Hopper**
  - Eight nodes
    - 16 cores, 2.4GHz AMD Opteron
  - Four nodes
    - 32 cores, 2.0GHz AMD Opteron
  - 128GB

- **Carver**
  - Four nodes
    - 8 cores, 2.66GHz Intel Nehalem
    - 48GB

Login Node Access

- **Connect (via ssh) to *load balancer***
  - edison.nersc.gov
  - hopper.nersc.gov
  - carver.nersc.gov
- **Load balancer selects login node based on:**
  - Number of connections
  - Memory of previous connections from same IP
    - If you login everyday, you'll probably end up on the same login node every time.

Login Node Usage

- **Login nodes are shared by many users, all the time**
- **Edit files, compile programs, submit batch jobs**
- *Some* **post-processing/data analysis**
  - IDL
  - MATLAB
  - NCL
  - python
- *Some* **file transfers**
  - Use data transfer nodes for large/long-running transfers
- **Please use discretion**
  - *All* users get frustrated by sluggish interactive response

Login Node Guidelines

- **Determine number of available cores**

  % grep processor /proc/cpuinfo | wc -l

- **Determine amount of physical memory**

  % grep MemTotal /proc/meminfo

- **Use "top" command to view process activity**

- **Limit use of parallel "make"**

  % make -j 4 all

- **Use *no more* than 50% of available cores**

- **Use *no more* than 25% of available memory**

- **NERSC will kill user processes if response becomes unacceptable**

Long-Term File Systems

- **Global home directories**
  - Source/object/executable files, batch scripts, input files, configuration files, batch job summaries (*not* for running jobs)
  - Backed up
  - 40GB permanent quota
  - $HOME

- **Global project directories**
  - Sharing data between people and/or systems
  - By PI request
  - Backed up if quota less than 5TB
  - 4TB default quota

Short-Term File Systems

- **Local scratch directories**
  - Cray (Edison, Hopper) only
  - Large, high-performance parallel Lustre file system
  - Not backed up; files purged after 12 weeks
  - Hopper: 5TB default quota; Edison: 10TB default quota
  - $SCRATCH, $SCRATCH1, $SCRATCH2, $SCRATCH3
- **Global scratch directories**
  - All systems
  - Large, high-performance parallel GPFS file system
  - Not backed up; files purged after 12 weeks
  - 20TB default quota
  - $GSCRATCH

File System Suggestions

- **Use $SCRATCH for running Hopper/Edison batch**

- **Use $GSCRATCH for running Carver batch**

- **Performance can be limited by metadata**
  - Do not store 1000s of files in single directory

- **Use "tar" to conserve inodes**

- **Use HPSS to archive important data**
  - Protection against hardware failure
  - Quota management

Shell Initialization Files

- **Standard dot files**
  - .bashrc, .profile, .cshrc, .login, etc.
  - Symbolic links to read-only files
    - Allows NERSC to provide common environment

- **Personal dot files**
  - Aliases, environment variables, modules, etc.
  - Use ".ext" files
  - .bashrc.ext, .profile.ext, .cshrc.ext, .login.ext, etc.

- **Use "fixdots" to start over**
  - Creates $HOME/KeepDots.<timestamp>
  - Restores all dot files to current default state

National Energy Research Scientific Computing Center